DS-GA 1008 - Deep Learning, Spring 2016

Assignment 1

Due: Thursday, Feb. 11th, 2016 at 4:55pm

---

1 Backpropagation

1. Warmup: Logistic regression is a pretty popular technique in machine learning to classify data into two categories. This technique builds over linear regression by using the same linear model but this is followed by the sigmoid function which converts the output of the linear model to a value between 0 and 1. This value can then be interpreted as a probability. This is usually represented as:

$$P(y = 1|x_{\text{in}}) = x_{\text{out}} = \sigma(x_{\text{in}}) = \frac{1}{1 + e^{-x_{\text{in}}}} \tag{1}$$

   where x_{in} as the name would suggest is the input scalar (which is also the output of linear model) and x_{out} is the output scalar.

   If the error backpropagated to x_{out} is $\frac{\partial E}{\partial x_{\text{out}}}$, write the expression for $\frac{\partial E}{\partial x_{\text{in}}}$ in terms of $\frac{\partial E}{\partial x_{\text{out}}}$.

2. Multinomial logistic regression is a generalization of logistic regression into multiple classes. The softmax expression is at the crux of this technique. After receiving n unconstrained values, the softmax expression normalizes these values to n values that all sum to 1. This can then be perceived as probabilities attributed to the various classes by a classifier. Your task here is to backpropagate error through this module. The softmax expression which indicates the probability of the i-th class is as follows:

$$P(y = i|X_{\text{in}}) = (X_{\text{out}})_i = \frac{e^{-\beta(X_{\text{in}})_i}}{\sum_k e^{-\beta(X_{\text{in}})_k}} \tag{2}$$

   What is the expression for $\frac{\partial (X_{\text{out}})_i}{\partial (X_{\text{in}})_j}$? (Hint: Answer differs when $i = j$ and $i \neq j$).

   The variables X_{in} and X_{out} aren't scalars but vectors. While X_{in} represents the n values input to the system, X_{out} represents the n probabilities output from the system. Therefore, the expression $(X_{\text{out}})_i$ represents the i-th element of X_{out}.

Figure 1: a DAG of modules

3. Figure 1 describes a simple network with the inputs being x_1 and x_2 (both scalars). Both inputs are passed through the sigmoid layer described above. Next, the outputs from the sigmoid layers are passed through max and min layers whose outputs can be represented as $o_{\min} = \min(i_1, i_2)$ and $o_{\max} = \max(i_1, i_2)$. The values output by the min and max layers are then summed in the final layer.

Suppose the error passed down (backpropagated) to y is $\frac{\partial E}{\partial y}$. Write the values for $\frac{\partial E}{\partial x_1}$ and $\frac{\partial E}{\partial x_2}$ in terms of $\frac{\partial E}{\partial y}$.

2 Torch (MNIST Handwritten Digit Recognition)

For this problem, you will need to run and understand the sample code located at `https://github.com/yjxiao/ds-ga-1008-a1`. The sample code is based on Clement Farabet's tutorial on supervised learning `http://code.madbits.com/wiki/doku.php?id=tutorial_supervised`. The difference is that we are using MNIST instead of SVHN. It is recommended that you walk through the tutorial in an itorch notebook.

You can clone and run the code on CIMS (use `https://www.cims.nyu.edu/webapps/content/systems/userservices/accounts/obtain` if you don't have a CIMS account) as follows:

```
you@crunchy1[Documents]$ git clone GIT_REPO_GIVEN_ABOVE
you@crunchy1[Documents]$ module load torch
you@crunchy1[Documents]$ cd ds-ga-1008-a1
you@crunchy1[ds-ga-1008-a1]$ th doall.lua
```

You need to save the predictions and submit them to the Kaggle competition. Once you have understood the starter code, try to improve the model and beat the competition. After you finalize your model, create a script `result.lua`, that generates a file named

`predictions.csv`. We will check that `predictions.csv` matches your final Kaggle submission. You will need to save your trained model as well to be able to generate the predictions via `result.lua`. The function `torch.save()` in the starter code does exactly this.

DO NOT submit your trained model file (the binary generated by torch during training) via email. Instead expose the file through CIMS. These files are generally huge. Use `http://cims.nyu.edu/webapps/content/systems/userservices/webhosting` for information on how to host on CIMS servers. Make sure you set the right permissions.

DO NOT train on the test data. Training on the test data is considered plagiarism. Note that the sample code does not use a validation set which we recommend you implement by yourselves.

Evaluation
Your grade for this section will be based on

- 30% - Kaggle performance. Full score as long as you beat the benchmark.

- 40% - Writeup on your model structure, training process, experiments, results, etc. We expect a rather formal report written with LaTeX.

- 30% - Simple, readable, commented code of final submitted script `result.lua` that is able to execute on the test data and generate a prediction file consistent with your final Kaggle submission.

Submission

Send your submission (writeup and `result.lua`) to your corresponding TA by the deadline. Merge your solutions to section 1 with the writeup from section 2. Include a link to the trained model file in the email. Please use the following title for your email.

`[DS-GA-1008 YOUR_TEAM_NAME] Submission A1`

If you are still forming a team, please introduce your team by sending an email to Jake `j.zhao@nyu.edu` with the following subject line and specify the team name, team members, and their NetIDs.

`[DS-GA-1008 YOUR_TEAM_NAME]`

We will be sending out Kaggle invitations as soon as you form a team. Please compete as teams on Kaggle with exactly the same team names.